

05-30-00

11



Express Mail Label Number:  
EL175466235US



**UTILITY PATENT APPLICATION TRANSMITTAL (37 CFR § 1.53(b))**  
**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

Commissioner of Patents and Trademarks  
Box Patent Application  
Washington, DC 20231

☒ Duplicate for fee processing

Sir: This is a request for filing a patent application under 37 CFR § 1.53(b) in the name of inventor(s):

**Dwight Sunada and Xuejun Yuan**

Title: **APPARATUS AND METHOD FOR FINDING THE LONGEST MATCHING  
PREFIX IN A ROUTER**

Application Elements:

- ☒ **23** Pages of Specification, Claims and Abstract
- ☒ **07** Sheets of Drawings
- ☒ **04** Pages Combined Declaration and Power of Attorney

Accompanying Application Parts:

- ☒ Assignment and Assignment Recordation Cover Sheet (recording fee of \$40.00 enclosed)
- ☐ Information Disclosure Statement with Form PTO-1449
  - ☐ Copies of IDS Citations
- ☐ Preliminary Amendment
- ☒ Return Receipt Postcard
- ☐ Other: **N/A**

Fee Calculations (37 CFR § 1.16):

	<u>No. Filed</u>		<u>No. Extra</u>	<u>Rate</u>	<u>Fee</u>
Basic Fee					\$ 690.00
Total Claims	30	-20 =	10	x \$18 =	\$ 180.00
Independent Claims	08	-03 =	05	x \$78 =	\$ 390.00
<input type="checkbox"/> Multiple Dependent Claim Presented					\$260= \$ xxx.xx
TOTAL FILING FEE:					\$1260.00

☒ Please charge Deposit Account no. 19-4974 in the amount of \$1260.00 to cover the total filing fees.

☒ The Commission is hereby authorized to charge any additional fees which may be required, including extension fees, or credit any overpayment to Deposit Account No. 19-4974.

☒ Please send correspondence to the following address:

**SUN MICROSYSTEMS**  
**901 San Antonio Road MS UPAL1-521**  
**Palo Alto, CA 94303**

Attention of Kenta Suzue Phone: (650) 336-7175; Fax (650) 336-0530.

Date: May 26, 2000



Kenta Suzue  
Registration No. 45,145

CERTIFICATE OF EXPRESS MAIL  
UNDER 37 C.F.R. § 1.10

"Express Mail" mailing label number: EL175466235US

DATE OF DEPOSIT: May 26, 2000

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. §1.10 on the date indicated above and is addressed to: Commissioner of Patents  
Box Patent Application  
Washington DC 20231



Kenta Suzue

**APPARATUS AND METHOD**  
**FOR FINDING THE LONGEST MATCHING PREFIX**  
**IN A ROUTER**

Dwight Sunada

Xuejun Yuan

## **FIELD**

This invention relates generally to routers, and more specifically to a method and apparatus for finding the longest matching prefix in a router.

## **BACKGROUND**

A router is a networking device that forwards an incoming packet to a targeted destination. The incoming packet includes a destination IP (Internet Protocol) address, which is matched against a list of entries in a table. Each entry includes an IP address prefix, a mask, and a port number. For every entry in the router table, the prefix is compared with the result of AND'ing the destination IP address with the mask. If an entry has a matching prefix, the entry supplies the port number, which indicates the next hop of the packet.

Alternatively, table entries can express the prefix and the mask in a combination called a trit representation. A trit includes an information state of "x" in addition to "1" and "0." The state of "x" indicates a state of "don't care." A trit representation may include several trits.

A specific example of a table follows:

<u>Prefix</u>	<u>Mask</u>	<u>trit representation</u>	<u>Port</u>
10.1.3.1	255.255.255.255	10.1.3.1	5
10.1.3.0	255.255.255.0	10.1.3.x	4
10.1.1.0	255.255.255.0	10.1.1.x	3
10.1.0.0	255.255.0.0	10.1.x.x	2
10.0.0.0	255.0.0.0	10.x.x.x	1

A destination IP address can match multiple entries of the table. For example, a destination IP address of 10.1.1.7 matches the bottom three entries in the above table. If the

table returns multiple matching entries for the destination address, the correct entry is the entry including the longest matching prefix. In the previous example, the longest matching prefix is 10.1.1.0 (the trit representation 10.1.1.x).

One goal of router design is to efficiently find the entry including the longest matching prefix from multiple matching entries. The IP address has 32 bits, and the IP address prefix in CIDR (Classless InterDomain Routing) for IPv4 (Internet Protocol Version 4) ranges in length anywhere from 1 bit to 32 bits. Thus, the destination IP address may match as as many as 32 entries of the table. A router should be able to determine the entry including the longest matching prefix from as many as 32 matching entries.

One approach is to construct the table such that the table includes the entries in order, for example, starting from entries including prefixes with 32 relevant bits (i.e., a mask of 255.255.255.255) and ending with entries including prefixes with 1 relevant bit. (i.e., a mask of 128.0.0.0). The entry including the longest matching prefix must then be the matching entry positioned earliest in the table.

One problem with this approach is that entries must be stored in predetermined positions. This leads to wasted memory space. For example, memory dedicated to entries with 32 relevant bits may be underutilized, and memory dedicated to entries with 31 relevant bits may be full; this represents a poor allocation of memory resources. Additionally, the process of updating entries maybe cumbersome. For example, if memory dedicated to entries with 31 bits is full, and an additional entry of a new prefix needs to be entered into the table, then memory space must be created elsewhere, and the memory allocated to the entries with 31 bits needs to be redefined, to keep intact the scheme of ordering entries.

Therefore, what is needed is a better solution to the problem of finding an entry including the longest matching prefix from multiple matching entries.

## **SUMMARY**

To overcome the limitations described above, and to overcome other limitations that will become apparent upon reading and understanding the present specification, disclosed is a system and method for finding the longest matching prefix in a router table for a destination IP address. In one embodiment, the system and method include a content addressable memory storing IP address prefixes, and an encoder coupled to the memory which finds the longest matching prefix. The IP address prefixes may be stored in the memory in a length independent manner. Each entry of an IP address prefix has a degree of relevance and either a matching status or a non-matching status. The system and method for reducing a number of multiple matching entries is achieved by changing a matching status of one or more less relevant entries in the multiple matching entries to a non-matching status. Less relevant matching entries that contain shorter IP address prefixes are automatically removed, leaving only the entry including the longest matching prefix.

## **DESCRIPTION OF THE DRAWINGS**

FIG. 1 is a block diagram showing a system including a length-based priority encoder;

FIG. 2 is a schematic diagram showing the length-based priority encoder of FIG. 1, including flag triggers and flags;

FIG. 3 is a logic diagram showing a flag trigger and a flag shown in FIG. 2;

FIG. 4 is a logic diagram showing another embodiment of a flag trigger, including find circuitry;

FIG. 5 is a logic diagram showing another embodiment of a flag trigger that locates the flag trigger away from other circuitry;

FIG. 6 is a logic diagram showing another embodiment of a flag; and

FIG. 7 is a logic diagram showing another embodiment of a flag trigger that includes different circuitry for receiving a signal to change a matching status to a non-matching status.

## **DETAILED DESCRIPTION**

This invention is described in a preferred embodiment in the following description with reference to the Figures, in which like numbers represent the same or similar elements. While this invention is described in terms of the best mode for achieving this invention's objectives, it will be appreciated by those skilled in the art that variations may be accomplished in view of those teachings without deviating from the spirit or scope of the invention.

FIG. 1 shows an exemplary system 100, which may be a router. The system 100 includes a CAM (content addressable memory) 110, an LPE (length based priority encoder) 120, and a memory 130, a control unit 140, and a group of entries 150. An entry of the group of entries 150 includes information found in the CAM 110, the LPE 120, and the memory 130. In one embodiment of system 100 operating as a router, the entry stores an IP address prefix and mask (or a combination of the prefix and mask known also as a trit representation) in the CAM 110; a code called an SLE (special length encoding) in the LPE 120; and a port number

associated with the IP address prefix in the memory 130. The CAM 110 stores the IP address prefixes and masks of the group of entries 150. These entries may be stored in a random order. For example, these entries may be stored in a length independent manner. A destination IP address of a packet is sent to the CAM 110. The CAM 110 compares the destination IP address with all of the prefixes and masks stored in the CAM 110 to find matching entries. Then the CAM 110 sends one or more signals indicating entries having a matching status to the LPE 120. The LPE 120 includes in the SLE information that signifies the degree of relevance of the entries having the matching status. The LPE 120 finds more relevant matches by removing less relevant matches from consideration, changing the matching status of the less relevant matches to a non-matching status. Then the LPE 120 sends one or more signals to the memory 130 indicating matching entries having greater relevance. In one embodiment, the matching entry having the greatest relevance is indicated to the memory 130. The output of the system 100 includes contents of the memory 130, such as a port number, corresponding to the match having the greatest relevance. The control unit 140 performs administrative functions.

An entry of the group of entries 150 has some degree of relevance depending on the number of "1"s in the mask of the entry. For example, given two entries:

a first entry: prefix 10.1.1.0, mask 255.255.255.0, trit representation 10.1.1.x, port 3

a second entry: prefix 10.1.0.0, mask 255.255.0.0, trit representation 10.1.x.x, port 2

the first entry has a greater degree of relevance and the second entry has a lesser degree of relevance, because the first entry has more relevant binary numbers in the prefix corresponding to more "1"s in the mask or fewer "don't care's" in the trit representation. The degree of relevance is also stored in the SLE of the LPE 120. Given a packet with a destination



IP address, such as 10.1.1.1, that results in both the first entry and the second entry of the preceding example sending a signal indicating a matching status to the LPE 120, the LPE 120 changes the matching status of the less relevant second entry to a non-matching status. As the only matching entry that remains, the first entry is found to include the longest matching prefix. The matching status of the first entry is communicated to the memory 130, and the port number 3 stored in the memory 130 is the output of the system 100.

In other embodiments, an entry of the group of entries 150 includes information, a degree of relevance characterizing the information, and output information associated with the information. An input is compared with the stored information. Signals are sent to indicate entries including information that match the input. After processing the degrees of relevance characterizing the matching entries, matching entries having lesser degrees of relevance become non-matching entries. The degree of relevance can reflect a number of significant digits, searching criteria for an input, or any other characterization of an entry. Finally, the system 100 communicates output information associated with remaining matching entries. The system 100 may be any device that finds a most relevant entry or most relevant entries for a given input.

In other embodiments, the LPE 120 reduces a number of matching entries to one or more matching entries of greater relevance.

In other embodiments, the LPE 120 reduces a number of matching entries to one or more matching entries of lesser relevance, for example, by identifying entries that should be removed.

FIG. 2 shows an exemplary LPE 120, which includes a plurality of trigger arrays 220 and a plurality of flags 230. The LPE 120 receives a plurality of match lines 210. Each of the plurality of match lines 210 carries either a matching signal or a non-matching signal from the



triggering status in a particular SLE are consecutive and appear on a rightmost side or a leftmost side. Other embodiments with nonconsecutive flag triggers having the triggering status and do not appear on the rightmost side or the leftmost side are also practical.

Assuming that a particular flag trigger has the triggering status and a particular flag has the non-triggered status, the flag trigger is associated with the flag if a signal sent from the flag trigger prompting the status of the flag to change from the non-triggered status to the triggered status causes the status of the flag to change from the non-triggered status to the triggered status. Similarly, assuming that a particular flag trigger has the triggering status and a particular flag has the non-triggered status, the flag trigger is not associated with the flag if a signal sent from the flag trigger does not cause the status of the flag to change from the non-triggered status to the triggered status.

A trigger array holding an SLE characterizing an entry having a greater degree of relevance includes flag triggers having the triggering status associated with at least every flag of the plurality of flags 230 associated by flag triggers included in a trigger array holding an SLE characterizing an entry having a lesser degree of relevance, and another flag. In the specific examples of SLEs above for one specific embodiment, an SLE characterizing an entry with fewer "x" values ("don't care"s) has "1" bits in every column that has a "1" bit included in an SLE characterizing an entry with more "x" values ("don't care"s). In one embodiment implementing the above SLE examples, a flag is associated with all flag triggers in the same column as the flag.

In other embodiments, the plurality of trigger arrays 220 may be any logical or physical structure that supports the plurality of flag triggers 240. In other embodiments, the plurality of

flag triggers 240 may be any logical or physical structure that holds values indicating whether or not the plurality of flags 230 should be triggered. In other embodiments, the plurality of flags 230 may be any logical or physical structure that holds a logical value indicating whether a triggering signal has been received from the plurality of flag triggers 240.

FIG. 3 shows one embodiment of a flag trigger having one of a triggering or non-triggering status. In one embodiment, 32 LPE cells 300 are included in each trigger array to comply with an IP address prefix that includes as many as 32 bits according to IPv4. In another embodiment, 128 LPE cells 300 are included in each trigger array to comply with IPv6 (Internet Protocol Version 6) which includes 128 bits in the IP address.

The LPE cell 300 uses a 2-phase clock, driven by a clock  $\phi 1$  and a clock  $\phi 2$ . A precharge phase occurs when the clock  $\phi 2$  is high and the clock  $\phi 1$  is low. During the precharge phase, a transistor M9 turns on and discharges both a gate of a transistor M11 and a source of a transistor M10. A signal  $\phi 2\_b$  is a complement of the clock  $\phi 2$ . The signal  $\phi 2\_b$  precharges an input line 310 of a single-ended sense amplifier X8 to a voltage  $V_{dd}/2$  and an output line 320 of a single-ended sense amplifier X8, to a voltage  $V_{dd}$ . The output line 320 acts as the flag in this embodiment. In this embodiment, the non-triggered status of the flag is a voltage at or near the precharge voltage  $V_{dd}$  of the output line 320, and the triggered status of the flag is a voltage near the ground voltage. Precharging the output line 320 causes a gate of a transistor M5 to be at a ground voltage. The signal  $\phi 2\_b$  turns on a transistor M0 and precharges a cancel line 330 to the voltage  $V_{dd}$ .

An evaluation phase occurs when the clock  $\phi 1$  is high and the clock  $\phi 2$  is low. An inverter X3 and an inverter X4 form an SRAM cell. The SRAM cell acts as the flag trigger in



ground voltage. A device X2 changes the status of the entry characterized by the current trigger array from the matching status to the non-matching status.

### Case 3

"Flag Trigger has Triggering Status and Flag has Triggered Status"

This case indicates that the current flag trigger caused or helped to cause the flag to change from the non-triggered status to the triggered status. In addition to the current flag trigger, another flag trigger having the triggering status and included in another trigger array associated with the flag may have helped change the current flag from the non-triggered status to the triggered status. But there is no indication that the matching entry characterized by the current trigger array has less relevance than the matching entry characterized by the different trigger array, and no reason to change the matching status of the current entry to the non-matching status. Transistor M5 remains off and the matching status carried by the match line 340 propagates through the device X2.

### Case 4

"Flag Trigger has Non-Triggering Status and Flag has Non-Triggered Status"

This case indicates that neither the current flag trigger nor any flag trigger associated with another trigger array characterizing another matching entry has prompted the current flag to change from the non-triggered status to the triggered status. Thus, there is no indication of an existence of another trigger array characterizing another matching entry having a greater relevance. Transistor M5 remains off and the matching status carried by the match line 340 propagates through the device X2.

Other circuit elements in FIG. 3 are a device X0 and a device X9. The device X0

forwards the remaining match line having a matching status on to the memory 130. The device X9 forwards a bitwise negated SLE to an extra level of circuitry. The device X9 allows aggregation of multiple blocks of circuitry to form a much larger table of entries.

FIG. 4 shows another embodiment of a flag trigger including circuitry for finding a particular trigger array from the plurality of trigger arrays 220. The find circuitry includes transistors M1, M3, M4, M7, and M8; and devices X1 and X6. The find circuitry allows the control unit 140 to delete an entry by helping the control unit 140 to find the entry to be deleted. To delete a particular entry, the control unit 140 submits the IP address prefix to the CAM 110 and the prefix's SLE to the LPE 120. The control unit 140 finds the entry to delete by checking for a match from both the CAM 110 and the LPE 120. Other embodiments take advantage of other voltages, other devices, and other terminals of devices.

FIG. 5 shows an embodiment that locates the flag trigger away from other circuitry. A bus 500 is "n" bits wide to deliver "n" flag trigger values.

FIG. 6 shows an embodiment replacing the single-ended sense amplifier X8, the input line 310, and the output line 320 with a transistor M12. In this embodiment, the flag is a drain of the transistor M12. This embodiment may be sufficient if there are few entries. An embodiment using the single-ended sense amplifier as the flag may be more appropriate when there are many entries, particularly due to practical speed considerations. Other embodiments take advantage of other voltages, other devices, and other terminals of devices.

FIG. 7 shows an embodiment that replaces the transistor M0 with a single-ended sense amplifier X10 and splits the the cancel line 330 into an input cancel line 710 and an output cancel line 720. During the precharge phase, the single-ended sense amplifier X10 precharges

the input cancel line 710 to the voltage  $V_{dd}/2$  and the output cancel line 720 to the voltage  $V_{dd}$ . During the evaluation phase, the transistor M5 starts to lower the voltage on the input cancel line 710. The single-ended sense amplifier X10 detects a slight voltage drop in the input cancel line 710 and quickly drives the output cancel line 720 to the ground voltage. Other embodiments take advantage of other voltages, other devices, and other terminals of devices.

Another embodiment uses a sense amplifier to sense a signal prompting a matching status to change to a non-matching status, when a trigger array includes many flag triggers. Many flag triggers mean that excessive capacitance along the cancel line 330 will slow down voltage decrease of the cancel line 330. An embodiment used for IPv6 that includes 128 flag triggers in the trigger array may respond more quickly with this embodiment.

While some embodiments of the present invention have been illustrated herein in detail, it should be apparent that modifications and adaptations to embodiments may occur to those skilled in the art without departing from the scope of the present invention as set forth in the following claims.



## CLAIMS

1. A router, comprising:  
  
a content addressable memory which stores Internet Protocol address prefixes in an order independent of lengths of the Internet Protocol address prefixes; and  
  
an encoder coupled to the content addressable memory which stores a plurality of codes, corresponding to the Internet Protocol address prefixes in the content addressable memory, and compares the codes corresponding to matching Internet Protocol address prefixes to find a longest matching Internet Protocol address prefix.
2. The router of claim 1, further comprising:  
  
a memory coupled to the encoder, the memory for storing a port number corresponding to each Internet Protocol address prefix in the content addressable memory and other information for routing an incoming Internet Protocol packet.
3. The router of claim 1, wherein the encoder includes circuitry for finding one of the plurality of codes.
4. The router for claim 1, wherein the encoder includes circuitry for deleting one of the plurality of codes.
5. The router of claim 1, wherein each of the plurality of codes indicates a number of relevant bits in the corresponding Internet Protocol address prefix.

6. The router of claim 5, wherein among the codes corresponding to matching Internet Protocol address prefixes, a code indicating a highest number of relevant bits indicates the longest matching Internet Protocol address prefix.

7. The router of claim 5, wherein the code indicates up to 32 relevant bits in the corresponding Internet Protocol address prefix.

8. The router of claim 5, wherein the code indicates up to 128 relevant bits in the corresponding Internet Protocol address prefix.

9. A method for finding a longest matching prefix for an Internet Protocol address, comprising:

storing Internet Protocol address prefixes in a content addressable memory in an order independent of lengths of the Internet Protocol address prefixes; and

comparing codes corresponding to matching Internet Protocol address prefixes in an encoder to find a longest matching Internet Protocol address prefix.

10. The method of claim 9, wherein the codes indicate numbers of relevant bits in the corresponding Internet Protocol address prefixes.

11. The method of claim 10, wherein among the codes corresponding to matching Internet

Protocol address prefixes, the code indicating a highest number of relevant bits indicates the longest matching Internet Protocol address prefix.

12. A method for generating a router table, comprising:
  - receiving an Internet Protocol address prefix;
  - storing the prefix in a router table entry;
  - generating a code indicating a number of relevant bits in the Internet Protocol address prefix; and
  - storing the code in the router table entry.
13. A system for reducing a number of multiple matching entries by changing a matching status of one or more less relevant entries in the multiple matching entries to a non-matching status, comprising:
  - a plurality of entries having one of the matching status and the non-matching status;
  - a plurality of flags having one of a triggered status and a non-triggered status; and
  - a plurality of trigger arrays coupled to the plurality of entries and the plurality of flags,wherein each of the plurality of entries is characterized by one of the plurality of trigger arrays.
14. The system of claim 13, wherein each of the plurality of trigger arrays includes:
  - a plurality of flag triggers having one of a triggering status and a non-triggering status.
15. The system of claim 14, wherein each of the plurality of flag triggers is associated with

one of the plurality of flags, and prompts the associated flag to change to the triggered status if the each of the plurality of flag triggers has the triggering status.

16. The system of claim 15, further comprising:

logic coupled to the plurality of trigger arrays and the plurality of flags, wherein the logic sends a signal prompting one of the plurality of entries to change to the non-matching status if

- i) an entry the plurality of entries is characterized by a trigger array of the plurality of trigger arrays including a flag trigger of the plurality of flag triggers having the non-triggering status and
- ii) the flag trigger of the plurality of flag triggers having the non-triggering status is associated with a flag of the plurality of flags having the triggered status.

17. The system of claim 16, wherein a first trigger array of the plurality of trigger arrays characterizing a first entry of the plurality of entries having a greater degree of relevance includes flag triggers of the plurality of flag triggers having the triggering status associated with at least

- a) every flag of the plurality of flags associated by flag triggers of the plurality of flag triggers having the triggering status included in a second trigger array of the plurality of trigger arrays characterizing a second entry of the plurality of entries having a lesser degree of relevance and
- b) another flag of the plurality of flags.

18. A circuit for selectively sending a signal to prompt an entry to have a non-matching status, comprising:

a flag having one of a triggered status and a non-triggered status;

a flag trigger coupled to the flag, wherein the flag trigger has one of a triggering status and a non-triggering status, and if the flag trigger has the triggering status, the flag trigger prompts the flag to have the triggered status; and

logic coupled to the flag trigger and the flag, wherein if the flag trigger has the non-triggering status and the flag has the triggered status, then the logic sends the signal to prompt the entry to have the non-matching status.

19. The circuit of claim 18, wherein the flag trigger includes an SRAM, the triggered status is a first binary value, and the non-triggered status is a second binary value.

20. The circuit of claim 18, wherein the flag includes an output of a single-ended sense amplifier, the triggered status is a first output voltage of the single-ended sense amplifier, and the non-triggered status is a second output voltage of the single-ended sense amplifier.

21. The circuit of claim 18, wherein the flag includes a terminal of a transistor, the triggered status is a first output voltage of the terminal of the transistor, and the non-triggered status is a second output voltage of the terminal of the transistor.

22. The circuit of claim 18, wherein the signal prompting the entry to have the non-matching status is sent to a single-ended sense amplifier.

23. The circuit of claim 18, further comprising:

a logic gate coupled to the logic, the logic gate receiving a first signal indicating one of the matching status and the non-matching status of the entry, and a second signal prompting the entry to have the non-matching status.

24. The circuit of claim 18, further comprising:

a find circuitry coupled to the flag trigger, wherein the find circuitry allows the entry to be found for deleting the entry.

25. A method of selectively prompting an entry to have a non-matching status, comprising:  
refreshing a flag to have a non-triggered status;

if a flag trigger associated with the flag has a triggering status, prompting the flag to have a triggered status; and

if the flag trigger associated with the flag has a non-triggering status and the flag has a triggered status, prompting the entry characterized by the flag trigger to have the non-matching status.

26. A method of operating a router, comprising:

receiving Internet Protocol address prefixes;

generating codes corresponding to a number of relevant bits in the Internet Protocol address prefix

receiving a packet with a destination Internet Protocol address;

comparing the destination Internet Protocol address to the Internet Protocol address

prefixes to find the Internet Protocol address prefixes that match the destination Internet Protocol address;

comparing the codes corresponding to the matching Internet Protocol address prefixes to find a longest matching Internet Protocol address prefix; and

sending the packet to a port corresponding to the longest matching Internet Protocol address prefix.

27. A method for reducing a number of multiple matching entries by changing a matching status of one or more less relevant entries in the multiple matching entries to a non-matching status, comprising:

defining a plurality of trigger arrays, wherein each trigger array of the plurality of trigger arrays includes a plurality of flag triggers;

characterizing each entry of a plurality of entries with a trigger array of the plurality of trigger arrays;

associating each flag trigger of the plurality of flag triggers with a flag of a plurality of flags;

sending a triggering signal from the flag trigger of the plurality of flag triggers to prompt a flag of the plurality of flags to change to a triggered status, if at least i) the flag trigger of the plurality of flag triggers and the flag of the plurality of flags are associated and ii) the flag trigger of the plurality of flag triggers has a triggering status; and

sending a signal prompting an entry of the plurality of entries to change to the non-matching status if i) the entry of the plurality of entries is characterized by a flag trigger of the

plurality of flag triggers having a non-triggering status, and ii) the flag trigger having the non-triggering status and the flag of the plurality of flags having the triggered status are associated.

28. The method of claim 27, wherein an entry of the plurality of entries having a greater degree of relevance is characterized by a trigger array including more flag triggers having the triggering status and an entry of the plurality of entries having lesser degree of relevance is characterized by a trigger array including fewer flag triggers having the triggering status.

29. The method of claim 28, further comprising:

after sending signals prompting entries of the plurality of entries to change to the non-matching status, determining an only entry having the matching status as a most relevant entry.

30. The method of claim 29, wherein the most relevant entry is an Internet Protocol address prefix determined to be the longest matching prefix.





100  
↓


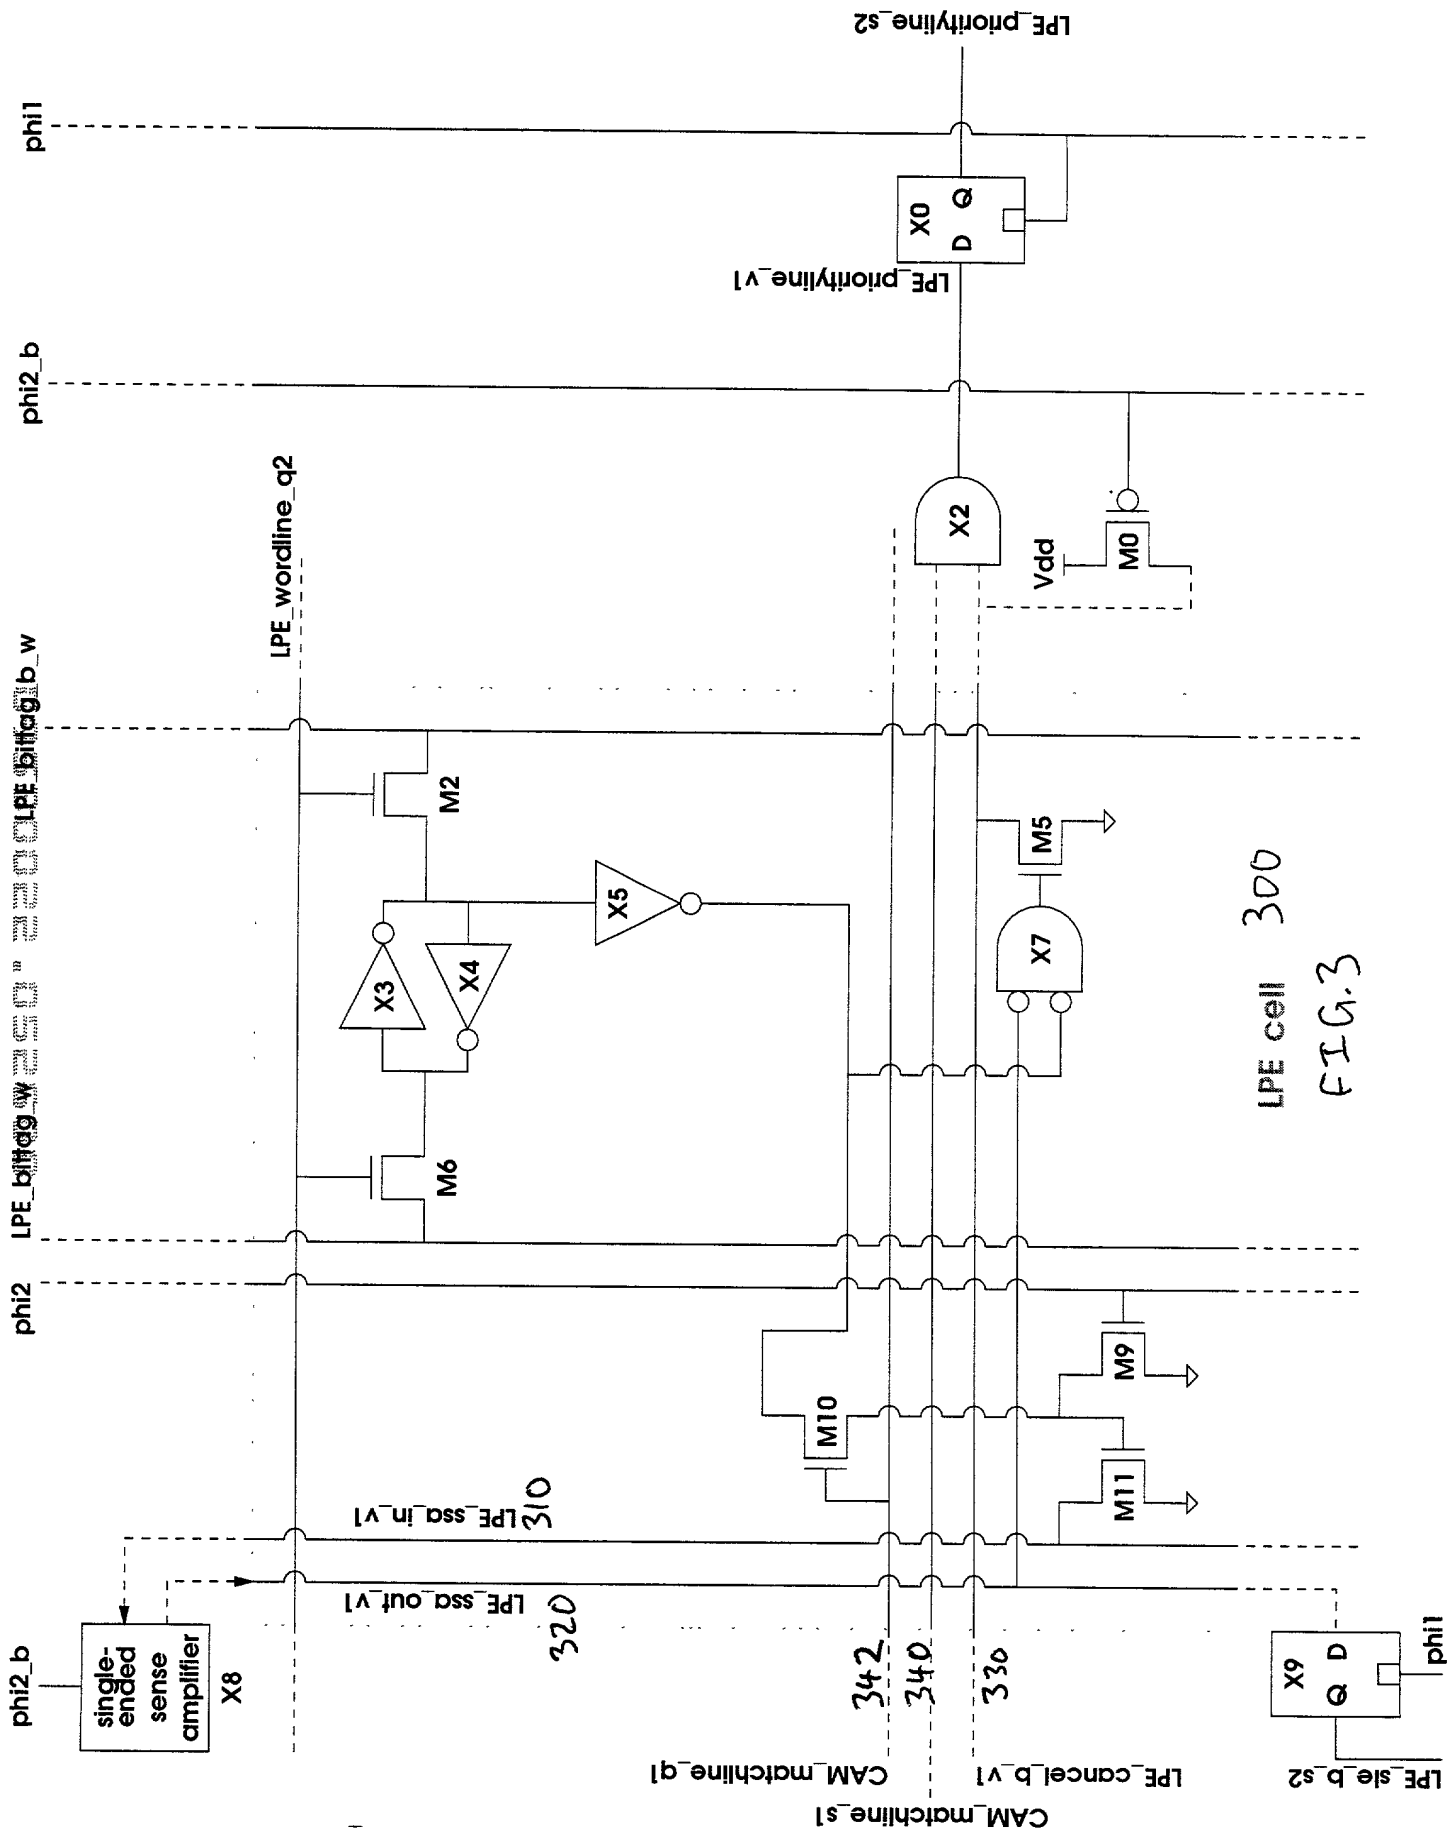
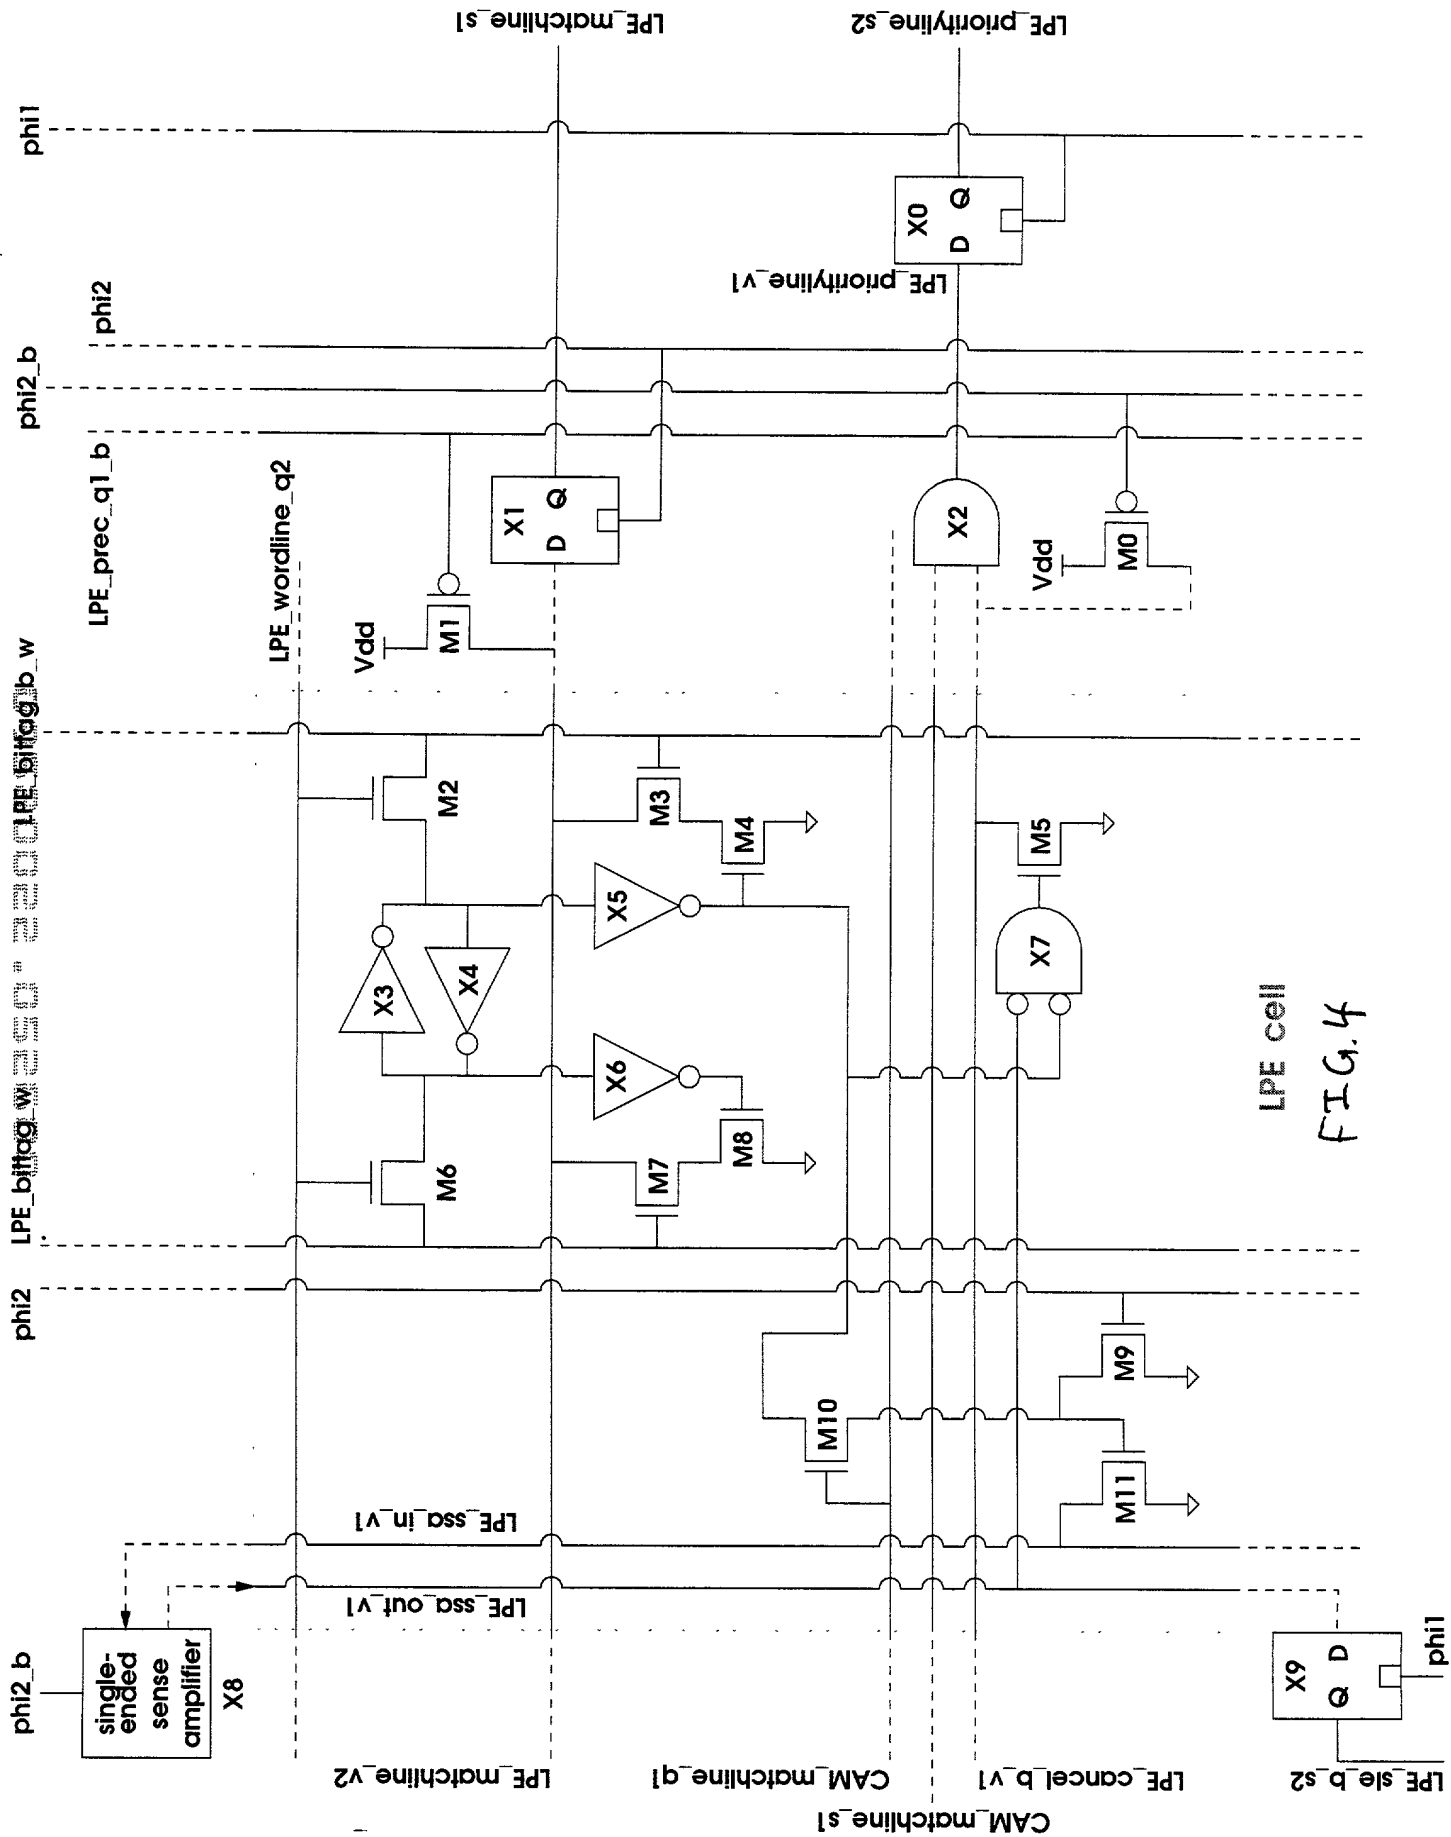


FIG. 1



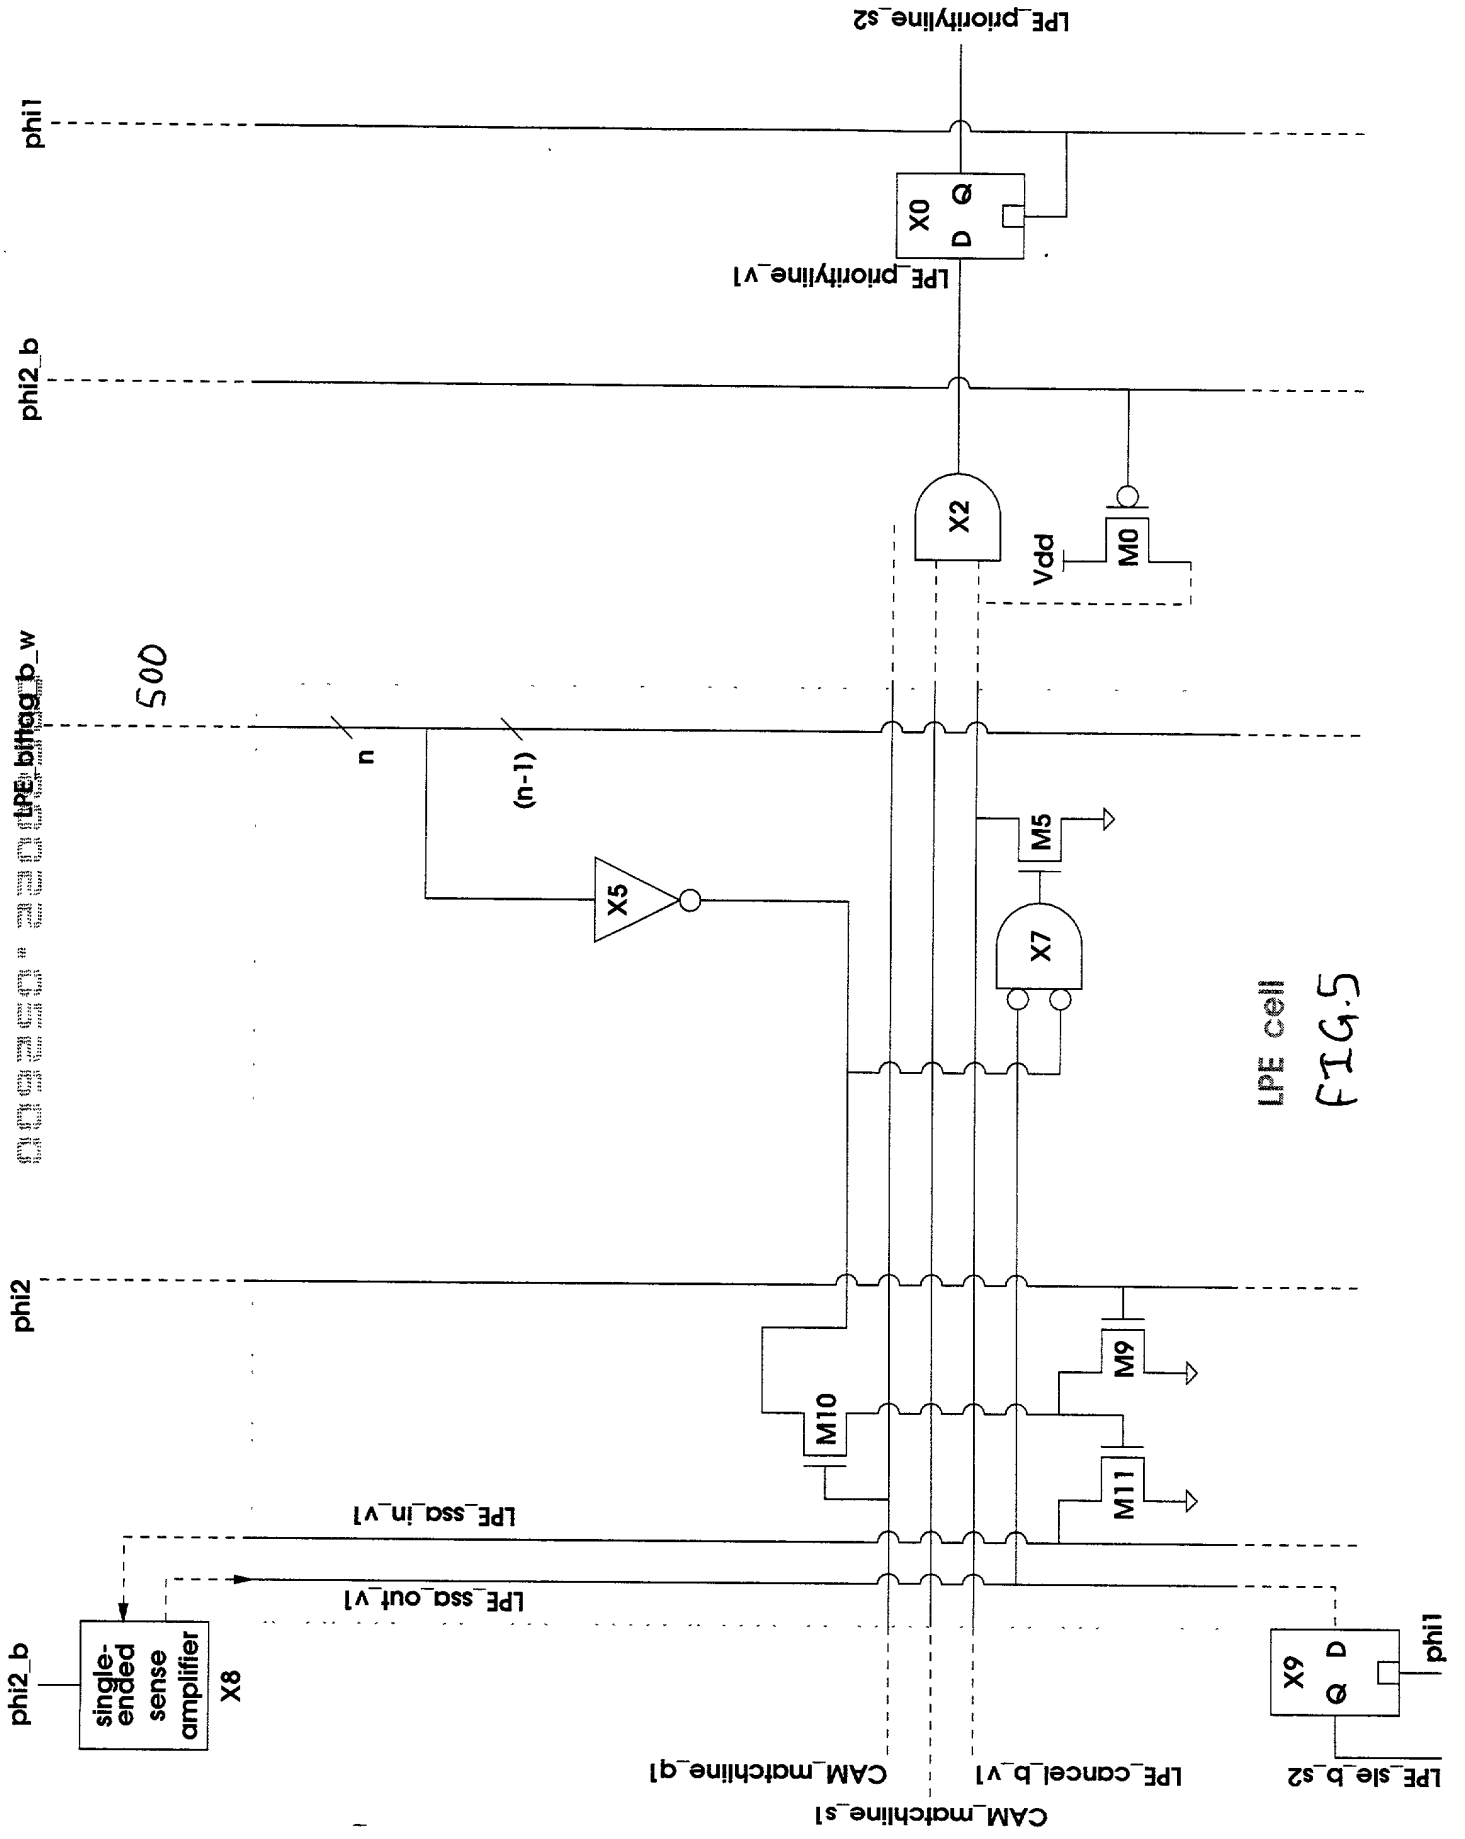


LPE cell 300  
FIG. 3

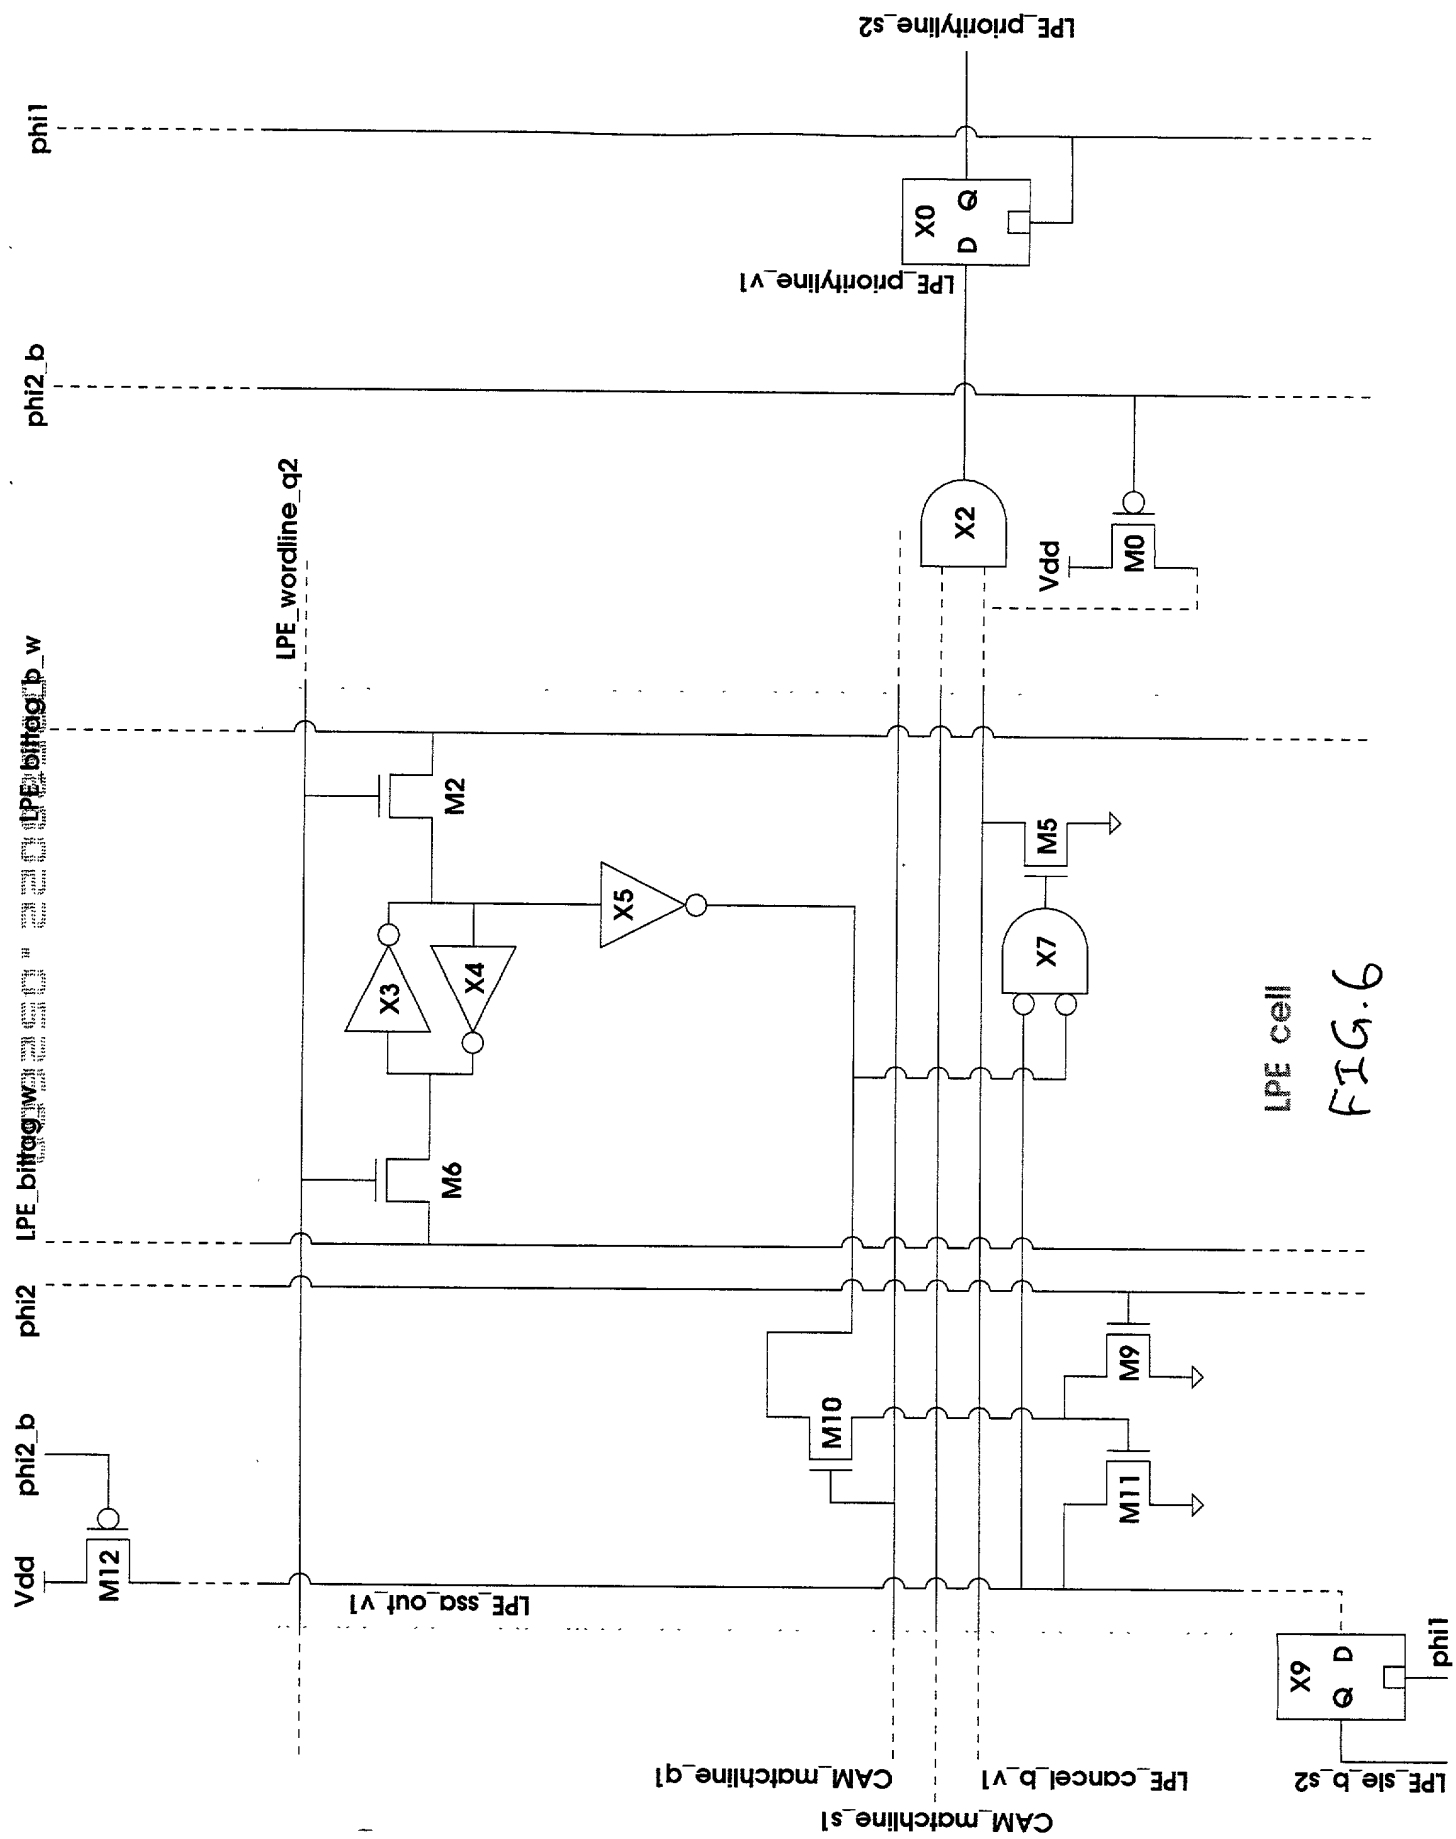


LPE cell

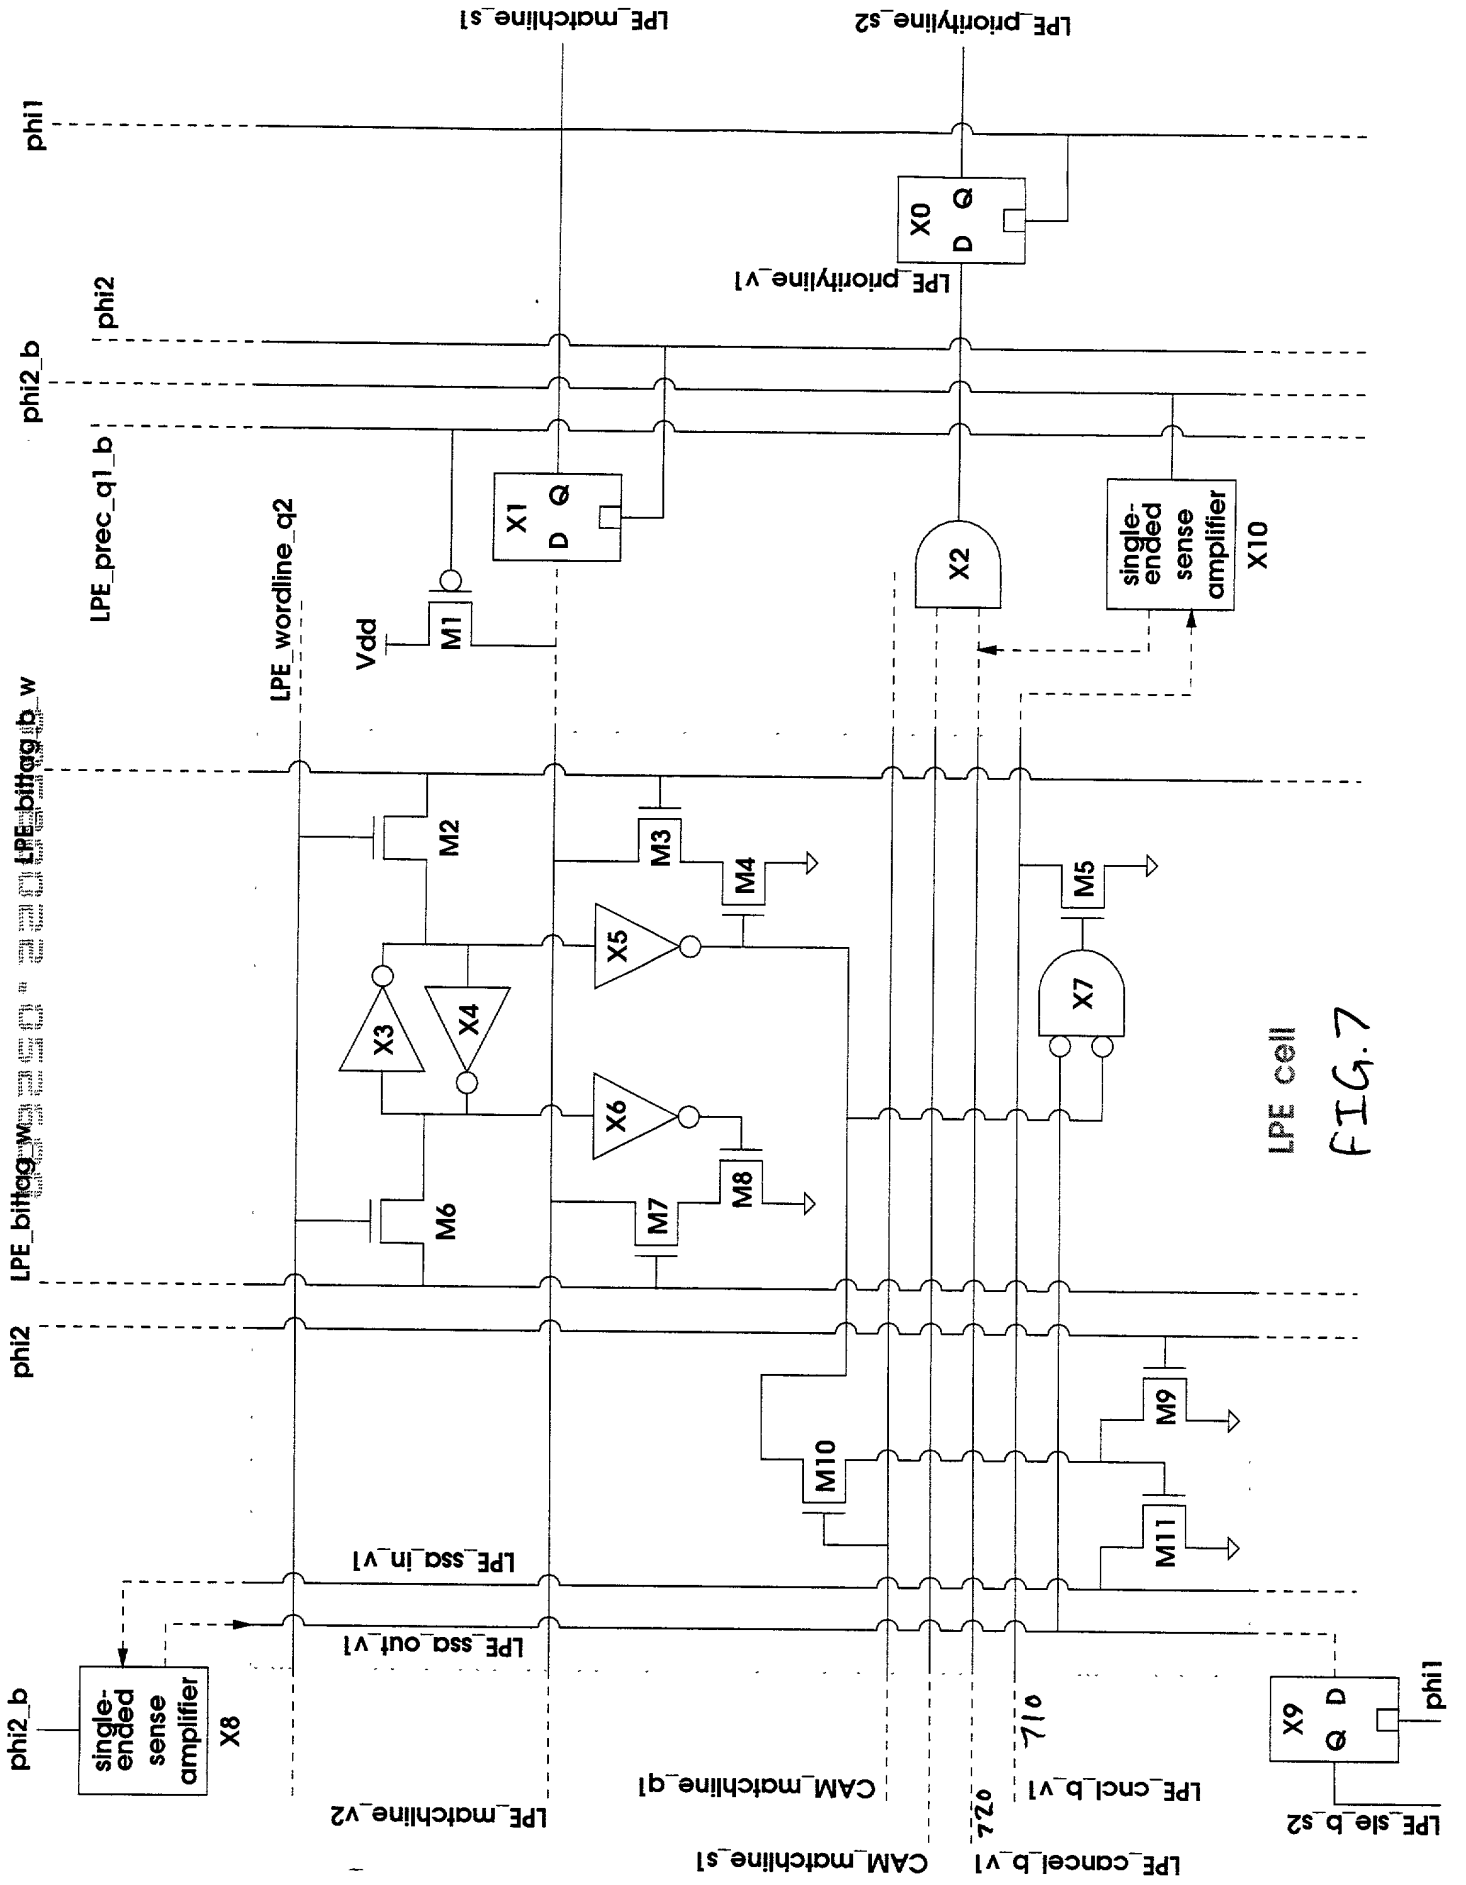
FIG.4



LPE cell  
FIG. 5



LPE cell  
FIG. 6



LPE cell  
FIG.7



**DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

**APPARATUS AND METHOD FOR FINDING THE LONGEST MATCHING  
PREFIX IN A ROUTER**

the specification of which

☒ is attached hereto.

☐ was filed on **xx/xx/xx** as United States Application Number **xx/xxx,xxx**  
or PCT International Application Number **N/A** and was amended on  
**N/A** (if applicable).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate or of any PCT international application having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)Priority Claimed

			YES	NO
xxxxxxxx.x (Number)	—	EP (Country)	xx/xx/xx (Date Filed)	<input type="checkbox"/> <input type="checkbox"/>
xxxxxxxx.x (Number)	—	JP (Country)	xx/xx/xx (Date Filed)	<input type="checkbox"/> <input type="checkbox"/>
N/A (Number)	—	N/A (Country)	N/A (Date Filed)	<input type="checkbox"/> <input type="checkbox"/>

I hereby claim the benefit under Title 35, United States Code, Section 119(e) of any United States provisional applications listed below

N/A (Application Number)	N/A (Filing Date)
N/A (Application Number)	N/A (Filing Date)

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s), or 365(c) of any PCT international application designating the United States of America, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

xx/xxx,xxx (Application Number)	xx/xx/xx (Filing Date)	xxxxxxxx (Status-Patent, Pending, Abandoned)
N/A (Application Number)	N/A (Filing Date)	N/A (Status-Patent, Pending, Abandoned)

POWER OF ATTORNEY

I hereby appoint

Kenneth Olsen, Reg. No. 26,493; Timothy J. Crean, Reg. No. 37,116; Joseph T. FitzGerald, Reg. No. 33,881; Robert S. Hauser, Reg. No. 37,847; Alexander E. Silverman, Reg. No. 37,940, Christine S. Lam, Reg. No. 37,489; Anirna Rakshpal Gupta, Reg. No. 38,275; Sean P. Lewis, Reg. No. 42,798; Michael J. Schallop, Reg. No. 44,319; Bernice B. Chen, Reg. No. 42,403; Kenta Suzue, Reg. No. 45,145; Noreen A. Krall, Reg. No. 39,734; Richard J. Lutton, Jr., Reg. No. 39,756; Monica D. Lee, Reg. No. 40,696 and Marc D. Foodman, Reg. No. 34,110 all of SUN MICROSYSTEMS, INC. with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Address all Correspondence to:

Sun Microsystems, Inc.  
901 San Antonio Road, MS UPAL1-521  
Palo Alto CA 94303

Attention of Kenta Suzue Phone: (650) 336-7175; Fax: (650) 336-0530.

SIGNATURES

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application, any patent issued thereon, or any patent to which this declaration is directed.

Full Name of Sole/First Inventor: DWIGHT SUNADA

Inventors Signature: *Dwight Sunada* Date 2000 MAY 26

Residence: Palo Alto, California Citizenship: USA

Post Office Address: 535 Everett Ave. Apt 605, Palo Alto, California 94301

Full Name of Second/Joint Inventor: XUEJUN YUAN

Inventors Signature: *Xuejun Yuan* Date: 5/26/00

Residence: San Jose, California Citizenship: People's Republic of China

Post Office Address: 6240 Goddess Court, San Jose, California 95129

Full Name of Second/Joint Inventor \_\_\_\_\_

Inventors Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence \_\_\_\_\_ Citizenship \_\_\_\_\_

Post Office Address \_\_\_\_\_

Full Name of Third/Joint Inventor \_\_\_\_\_

Inventors Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence \_\_\_\_\_ Citizenship \_\_\_\_\_

Post Office Address \_\_\_\_\_

Full Name of Fourth/Joint Inventor \_\_\_\_\_

Inventors Signature \_\_\_\_\_ Date \_\_\_\_\_

Residence \_\_\_\_\_ Citizenship \_\_\_\_\_

Post Office Address \_\_\_\_\_